

Container solutions for HPC Systems: A Case Study of Using Shifter on Blue Waters

Maxim Belkin University of Illinois at Urbana-Champaign National Center for Supercomputing Applications Urbana, IL, USA mbelkin@illinois.edu

Hon Wai Leong National Center for Supercomputing Applications University of Illinois at Urbana-Champaign Urbana, IL, USA hwleong@illinois.edu Roland Haas University of Illinois at Urbana-Champaign National Center for Supercomputing Applications Urbana, IL, USA rhaas@illinois.edu

Eliu A. Huerta National Center for Supercomputing Applications Department of Astronomy University of Illinois at Urbana-Champaign Urbana, IL, USA elihu@illinois.edu

> Mark Neubauer Department of Physics University of Illinois at Urbana-Champaign Urbana, IL, USA msn@illinois.edu

Galen Wesley Arnold University of Illinois at Urbana-Champaign National Center for Supercomputing Applications Urbana, IL, USA gwarnold@illinois.edu

> David Lesny Department of Physics University of Illinois at Urbana-Champaign Urbana, IL, USA ddl@illinois.edu

ABSTRACT

Software container solutions have revolutionized application development approaches by enabling lightweight platform abstractions within the so-called "containers." Several solutions are being actively developed in attempts to bring the benefits of containers to high-performance computing systems with their stringent security demands on the one hand and fundamental resource sharing requirements on the other.

In this paper, we discuss the benefits and short-comings of such solutions when deployed on real HPC systems and applied to production scientific applications. We highlight use cases that are either enabled by or significantly benefit from such solutions. We discuss the efforts by HPC system administrators and support staff to support users of these type of workloads on HPC systems not initially designed with these workloads in mind focusing on NCSA's Blue Waters system.

PEARC '18, July 22-26, 2018, Pittsburgh, PA, USA

© 2018 Copyright held by the owner/author(s). Publication rights licensed to the Association for Computing Machinery.

ACM ISBN 978-1-4503-6446-1/18/07...\$15.00

https://doi.org/10.1145/3219104.3219145

CCS CONCEPTS

• Computing methodologies \rightarrow Massively parallel and highperformance simulations; • Software and its engineering \rightarrow Virtual machines; • Applied computing \rightarrow Astronomy; Physics;

KEYWORDS

Petascale, Reproducibility, Data Science

ACM Reference Format:

Maxim Belkin, Roland Haas, Galen Wesley Arnold, Hon Wai Leong, Eliu A. Huerta, David Lesny, and Mark Neubauer. 2018. Container solutions for HPC Systems: A Case Study of Using Shifter on Blue Waters. In *PEARC '18: Practice and Experience in Advanced Research Computing, July 22–26, 2018, Pittsburgh, PA, USA.* ACM, New York, NY, USA, 8 pages. https://doi.org/10. 1145/3219104.3219145

THE RISE OF CONTAINERS

The enormous growth of computing resources has forever changed the landscape and pathways of modern science by equipping researchers with the apparatus that is impossible to realize experimentally. The great examples are data- and compute-enabled machine and deep learning algorithms that control self-driving cars; precise *in silico* studies of complete virus capsids that further our understanding of their pathogenic pathways; and the fascinating studies of gravitational waves that resonate around our Universe.

This growth of computing resources has been multi-directional: they increased in their availability, performance, and assortment. A typical computer today is equipped with Graphics Processing Units

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

(sometimes combined with Central Processing Units), abundant Random Access Memory, hard drives that can store Tera bytes of data, and many other, often highly specialized, hardware. Supercomputers, the high-performance computing (HPC) resources that drive modern science, have additional levels of complication with their stringent security demands, fundamental resource sharing requirements, and many specialized libraries that enable use of the underlying hardware at its peak performance.

Variations in hardware and software stacks across leadershipclass computing facilities have raised a great deal of concern among researchers with the most prominent one being *reproducibility* of computational studies. To ensure reproducibility, it is critical to use portable software stacks that can be seamlessly deployed on different computing facilities with their specific architectures. This need has driven the development of software solutions that abstract the underlying hardware away from the software. Today's most popular examples include container solutions like Docker, virtual machine solutions such as VirtualBox and VMWare Workstation, and others. These solutions differ in levels at which abstractions take place (hardware, OS, *etc.*), abstracted and required resources, as well as all auxiliary tools that together comprise their *ecosystems*. In this article we focus on a container solution for HPC systems: *Shifter*.

In addition to facilitating the use of complex software stacks within the HPC community, containers have also played a central role in a new wave of innovation that has fused HPC with highthroughput computing (HTC)—a computing environment that delivers a large amount of computing power over extended periods of time. A number of large scientific collaborations have made use of containers to run computationally demanding HTC-type workflows using HPC resources.

In this paper, we showcase a number of efforts that have successfully harnessed the unique computing capabilities of the Blue Waters supercomputer, the NSF-supported, leadership-class supercomputer operated by National Center for Supercomputing Applications (NCSA), to enable scientific discovery. We focus on efforts that have been spearheaded by researchers at NCSA and the University of Illinois at Urbana-Champaign [7, 17]. These efforts provide just a glimpse of the wide spectrum of applications in which containers help advance fundamental science: from the discovery of gravitational waves from the collision of two neutron stars with the LIGO and Virgo detectors [1], to the study of the fundamental building blocks of nature and their interactions with CERN's Large Hadron Collider (LHC).

SHIFTER ON BLUE WATERS

Shifter [3] is a container solution that is designed specifically for HPC systems and which enables hardware abstraction at the OS level. Figure 1 illustrates the workflow of a typical *Shifter* v16.08.3 job on the Blue Waters supercomputer.

It is extremely easy to get started using *Shifter* as all one has to do is provide an additional generic resource request: -1 gres=shifter16, either on the command line or in a job batch script. This PBS directive is mandatory for any *Shifter* job on Blue Waters as it instructs



Figure 1: Architecture of Shifter implementation on Blue Waters

system's workload manager to execute special *prologue* and *epilogue* scripts before and after job execution in order to set up and tear down the container environment on all compute nodes.

The other immediate advantage of *Shifter* is that it can work with Docker images—one of the most popular container formats—out of the box. There are two ways to specify which container to use in a batch job: either as a PBS directive -v UDI=<image:tag> or as an argument --image=<image:tag> to the shifter command provided by *Shifter*. In both cases, image corresponds to repository/imagename on Docker Hub.

When UDI is specified as a PBS directive, the prologue script communicates with the *Shifter* image gateway to check that the image exists and download it if it does not. The image gateway then applies site-specific environment changes to the image and converts it into a *squashfs*-formatted image file, typically referred to as User Defined Image, or UDI. The prologue script then proceeds to mount the UDI on all compute nodes allocated to the job. The UDI image file is stored on the Lustre file system and subsequent jobs requesting the same image can use it without repeating all of the above steps. Upon completion of such as job, the epilogue script unmounts the UDI from the compute nodes and performs sitespecific procedures necessary to properly clean up the environment on compute nodes. Both, site-specific environment changes to the downloaded image and cleanup procedures are specified by the system administrators.

The alternative way to specify UDI is by supplying it as an argument to the --image flag of the shifter command. In combination with Blue Waters' Application Level Placement Scheduler (ALPS) task launcher, to run an application within a container environment one can use the following command:

\$ aprun -b -- shifter --image=<image:tag> -<application>

The shifter command above initiates a series of operations which are similar to those executed by the *prologue* script of the workload manager. However, it not only provides the option to select container environment "on-the-fly," but also allows *Shifter* users to use several different images within the same job! This method is arguably best suited for single-task applications. For containerized MPI applications it is still recommended to use PBS directives to set up the container environment on all compute nodes. The core image gateway manager of *Shifter* is designed as a RESTful service. It is written in Python language and depends on multiple software components:

Flask - A Python-based framework that provides a RESTful API as an interfacing layer between user requests and the underlying image gateway. The use of RESTful API replaces local Docker engine as a gateway for users to request containers from a Docker registry.

MongoDB - A distributed database to store metadata of available container images and their operational status: whether the image is still being downloaded, its conversion status, its readiness to use, or any failure encountered.

Celery - A Python-based asynchronous and distributed task queueing system to service user requests. *Celery* provides better scalability for multiple requests through queueing and dispatch to a distributed pool of workers.

Munge - An authentication service for creating and validating credentials, designed to be highly-scalable, which is ideal for highperformance computing environment. *Shifter* uses *Munge* to authenticate user requests from clients to the gateway manager.

Redis - An in-memory data structure store, used as a database, cache and message broker to support *Celery*'s functionality. It captures the operational state of the *Shifter* image gateway service to allow live reconfiguration of *Shifter* (service restart) without interrupting any current operations.

When compared to its previous version [8], *Shifter v*16.08.3 features improved functionality and performance. Yet, just like the predecessor, it still introduces a noticeable overhead for system administrators who are responsible for its back-end operation because it relies on a number of very different components working seamlessly and with no interruptions. Without a doubt, it is much harder to troubleshoot an issue that involves *Shifter* as its root cause may not come from the tool itself but from one of its dependencies.

During the production use of *Shifter* on Blue Waters, the following issues have been encountered:

1. *Stale "PENDING" state*. When downloading containers from Docker registry, the status would stay in "PENDING" state indefinitely until its metadata is manually deleted from MongoDB's database. This usually happens when a user aborts the download of a large container from the Docker registry before the download completes.

2. *False "READY" state.* Status of a container image would indicate "READY", even though *Shifter* has, in fact, failed to mount the UDI on the compute nodes due to the unfinished download of the employed container image. The troublesome UDI file has to be removed from the storage and Docker image has to be re-downloaded.

3. Persistent out-of-memory issue on gateway host. There was an incident when the gateway manager caused the gateway host to run out of memory and, consequently, go down because multiple threads were downloading the same image from the Docker registry. Upon rebooting the host and restarting all of the services required by *Shifter*, multiple threads resumed their downloads leading to repeated failures. Subsequent restarts did not produce expected results. The solution was to remove the Redis dump DB file.

4. Failure to mount UDI when Munge is not running. Munge is crucial for *Shifter* to function properly. A compute node that does

not have Munge service running would not be able to authenticate with the *Shifter* image gateway and thus would fail to mount a UDI.

5. Failing to run at scale. The major challenge that we had to address on Blue Waters was to make *Shifter* jobs run at scale. The issue was caused by the bottleneck in getgrouplist and getgrgid functions that *Shifter* uses to set up the containers on compute nodes. These two functions query local passwd and group files and LDAP. Because Blue Waters does not store regular user and group information in passwd and group files, *Shifter* was trying to get the gids of the executing user from LDAP. For jobs with a large node count this step results in a large number of concurrent requests being sent to the underlying LDAP server. As a result, not all requests receive a response from the server. To work around this issue, we had to turn on the Name Service Cache Daemon (NSCD) service on all compute nodes allocated to *Shifter* jobs. The NSCD service caches LDAP entries on the compute nodes and, therefore, enables their fast lookup.

MPI APPLICATIONS IN SHIFTER JOBS

MPI is a performance-critical component of and *de facto* the standard for writing applications that run at scale. Therefore, for systems like Blue Waters it is crucial to understand the overhead that applications within *Shifter* UDIs have to pay in order to run on multiple nodes. To estimate this overhead, we compared *Shifter* to the Cray Linux Environment (CLE) using the OSU Micro-Benchmarks.

A selection of representative benchmarks were run: MPI_BCaSt, MPI_Reduce, MPI_AlltoAll, and MPI_AlltoAllv. Tests were performed on 64 and 1024 ranks, that correspond to 4 and 64 compute nodes on Blue Waters, correspondingly, see Figure 2. Employed *Shifter* image was based on clean Centos 7 Docker image, with MPICH v3.2 and OSU Micro-Benchmarks v5.3.2 installed from source. Our results suggest that MPI performance in CLE and *Shifter* is statistically the same. This stunning result is not surprising, however, because *Shifter* is able to use the Cray MPI low level communication libraries through the MPICH ABI compatibility initiative.

We set up MPI benchmarks in a way that made *Shifter* the only component that could significantly affect the results. In particular, the binaries were built with tools provided by the GNU Programming environment on Blue Waters (*PrgEnv-gnu*) for the CLE tests, and with mpicc that calls GNU compilers in the *Shifter* UDI that was based on Centos7 Docker image. Tests were run from the same batch job, minimizing the effect of node placement and Gemini network paths on the obtained results as much as possible. Only the variable network traffic that is associated with the production machine and that we don't have control over could have impacted the results. Because the results were obtained from the same jobs, we're confident that they are valid and reproducible.

I/O PERFORMANCE IN SHIFTER JOBS

Performance of read and write operations is crucial for the HTC type of applications that deal with lots of data. To see if *Shifter* imposes any input/output (I/O) overhead, we ran the IOR MPI I/O benchmark (https://github.com/hpc/ior, commit aa604c1) using 16 nodes and 7 cores for reading and writing operations on each node. Blue Waters runs the IOR benchmark on a regular schedule using the Jenkins testing infrastructure. To make the comparison



Figure 2: Comparison of OSU micro benchmarks' results measuring MPI performance in *Shifter* and Cray's native Linux Environment on Blue Waters using 64 compute nodes and 1,024 MPI Ranks. (a) MPI_Alltoall and MPI_Alltoallv. (b) MPI_Bcast and MPI_Reduce.

between the tests meaningful, we used the same input and node layout in our *Shifter* tests. Our results suggest, that there is no substantial differences between I/O performance in the native Cray Linux Environment (Jenkins test case) and the *Shifter* case, see Figure 3.

START-UP TIME OF SHIFTER JOBS

Shifter enables many new types of applications take advantage of HPC resources. As such, one may expect untraditional for HPC usage patterns to emerge. For example, starting production simulations or different stages of analysis from within a *Shifter* image multiple times within a job. To help users with such applications better utilize HPC resources, we analyzed the start-up time of *Shifter* jobs for User-Defined Images of two sizes: 36 MB and 1.7 GB. The results are shown in Figure 4.

We investigated how start-up time of a *Shifter* job depends on the number of nodes used by the job that exploits only 1 processor on each node. In our tests, we started *Shifter* jobs in two different ways: 1. by specifying UDI at the time the job was submitted, and 2. by specifying UDI as an argument to the shifter command. All of our tests suggest that start-up time of a *Shifter* job is practically independent of the size of the User-Defined Image! However, we find that for jobs using less than 256 nodes, the dependence of



Figure 3: Comparison of IOR benchmark results of IO performance in *Shifter* and Cray's native Linux Environment on Blue Waters using 16 compute nodes and 7 cores for reading and writing operations on each node.

the start-up time on node count is **sublinear**, beyond 256 nodes the dependence becomes linear, and beyond 2,048 – superlinear, see Figure 4 a.

We also studied the dependence of *Shifter* job start-up time on the number of MPI processes used on each node. All of these tests were performed using 80 compute nodes. And again, we find startup time to be practically independent of the size of the User-Defined Image we use. However, we find that when we specify UDI at the time we submitted the job, aprun calls take the same amount of time regardless of the number of processes on each node we request. This behavior is opposite of what we observe when we specify UDI as an argument to the shifter command. This observation suggests that if multiple calls to applications within the same UDI are necessary in a single job, it is advisable to specify UDI at the time the job is submitted.

CODES USING SHIFTER ON BLUE WATERS

Shifter was added to Blue Waters system in September of 2016 and was first used in a production simulation in January of 2017 by the ATLAS project to analyze data from the CERN's Large Hadron Collider [11]. The science team worked with the Blue Waters project to set up and test *Shifter*. The tested version of *Shifter* was then officially presented in a monthly user call in February of 2017 [2].

In order to learn about the codes that benefit from *Shifter* on Blue Waters, we collected information about its usage by analysing accounting records for the period from September, 2016 to March, 2017. In our analysis we did not include the simulations that ran for less than 1 hour. Interestingly enough, however, we found no significant difference in the distribution of codes when using a 5 minute Container solutions for HPC systems



Figure 4: Start-up time of *Shifter* jobs on Blue Waters. (a) Dependence of a *Shifter* job start-up time on the number of nodes. Start-up time is found to be practically independent of the way we specify which UDI to use in a job and the size of that image. (b) Dependence of the start-up time of a *Shifter* job that uses 80 nodes on the number of MPI processes used on each node. When UDI is specified at the time the job is submitted, job start-up time does not change when with the number of MPI processes used on each node!

"cutoff" instead. Figure 5 shows the distribution of node-hours consumed by different codes during the analyzed period. As is clear from Figure 5 the majority of the node-hours used with Shifter were consumed by ATLAS, NANOGrav, and LIGO projects. All of them are HTC codes that employ a large number of short and independent tasks that represent a trivially parallelizable workload. On HPC systems like Blue Waters, such codes typically use the so-called "pilot jobs" [9] that reserve compute nodes and aggregate them to a large shared compute pool of the HTC workflow manager. All three codes employ HTCondor [16] as the workflow manager and scale well to a large number of nodes. This scalability is achieved by using multiple pilot jobs to allow the workflow manager to release compute nodes when there are not enough tasks to utilize all provided resources. PySCF, QWalk, and QuantumEspresso represent "traditional" MPI-based HPC codes that utilize all allocated compute nodes and, therefore, benefit from the Shifter's ability to support MPI from within the containers. Finally, PowerGRID [4] is a modern, multi-GPU MPI code for reconstructing images obtained with the Magnetic Resonance Imaging technique. Figure 6 shows the distribution of node-hours used each month among the codes. As can be seen from Figure 6, Shifter has not been used continuously by any single code or science group on Blue Waters.

PEARC '18, July 22-26, 2018, Pittsburgh, PA, USA



Figure 5: Thousands of node-hours consumed by different codes using *Shifter* in the period 2016/09 – 2017/03. Of the codes shown, ATLAS, NANOGrav, and LIGO are well established high throughput computing workflows. PySCF, QWalk and QuantumEspresso are traditional HPC codes that employ MPI to achieve parallelization. PowerGRID is a GPU-enabled MPI code that can employ multiple GPUs on the compute nodes.



Figure 6: Node-hours used by *Shifter*-enabled codes on Blue Waters since 2016. The three early adopters—ATLAS, NANOGrav, and LIGO—employ typical for HTC workflows with multiple pilot jobs and HTCondor serving as the main workflow mananger. PySCF, QWalk and QuantumEspresso are traditional for HPC computational physics and chemistry codes. PowerGRID is a new MPI- and GPU-enabled code for MRI image reconstruction.

For codes such as NANOGrav and LIGO, this is due to the nature of their discrete analysis "Campaigns" during which collected data is analysed. Codes such as PowerGRID are still in the early stages of exploring the capabilities of *Shifter*. A follow-up study is necessary to determine if the observed non-continuous usage pattern is typical for applications that use *Shifter*.

Finally, Table 1 shows the number of nodes used by different applications on Blue Waters. As one can see from Table 1, most *Shifter*

Table 1: Top science applications and projects that use *Shifter* on Blue Waters. Columns show the number of nodes used in a typical job (Nodes), number of jobs ran (Frequency), and the total charge for the jobs (Node-Hours). The top three science projects that consumed the most resources while using *Shifter* are LIGO, ATLAS, and NANOGrav.

Code	Nodes	Frequency	Node-Hours
LIGO	10	8	1, 990
LIGO	50	5	2,650
LIGO	100	1	1,070
LIGO	500	1	6,030
LIGO	5,000	4	127,000
ATLAS	16	311	115,000
NANOGrav	1	1, 485	28, 500
NANOGrav	100	2	3,010
PowerGRID	800	1	4,870
PySCF	1	419	10,000
QWalk	1	1, 138	13, 200
QuantumEspresso	2	422	2,910

jobs are small (16 nodes or less) with only LIGO and PowerGRID attempting to scale up to larger node counts. This can be understood considering that available HTC tasks may not be sufficient to keep thousands of cores busy. Yet, an HPC system can not release just a fraction of nodes that are part of a job. This is the main reason for using multiple pilot jobs that can be terminated when necessary. The optimal size and number of pilot jobs depends on multiple factors such as the latency of the HPC scheduler, the length of each task, the backlog of available tasks in the workflow manager, and the "cost" of having idle nodes. Therefore, exploratory HTC runs use many small pilot jobs to determine the optimum quantities while only a few large pilot jobs are then used for production simulations, analysis, and testing.

ATLAS, NANOGrav, and LIGO

A lion's share of node-hours consumed by *Shifter* jobs on Blue Waters is associated with the three big state-of-the-art research projects: ATLAS, NANOGrav, and LIGO. Availability of sufficient computing resources was crucial for their Nobel prize-winning works that detected the Higgs boson and gravitational waves in 2013 and 2017, respectively.

Because all three codes use an Open Science Grid (OSG) [12]derived workflow, the challenges they face and behaviour they exhibit are very similar. Figure 7 shows a typical setup when using Blue Waters and *Shifter* as a compute resource in the OSG. For simplicity, we use LIGO as a stand-in for all three codes but the setup is, essentially, identical for all three projects.

The LIGO Scientific Collaboration employs HTCondor to analyse the data recorded by the LIGO detector, requiring that data from a repository in Nebraska, USA is transferred to a computing facility for processing. *Shifter* enabled LIGO collaboration to use an OSGready Docker image on Blue Waters, eliminating the need to adapt the image for each resource provider. This allowed LIGO to use an



Figure 7: Interaction between an science project data repository, the Open Science Grid and Blue Waters. Import/Export (IE) nodes are Blue Waters' dedicated nodes that are used for file transfer. Figure reproduced from [7].

operating system environment which is certified by the collaboration for a detection campaign and that matches the environment found on OSG resource providers: CentOS instead of Blue Waters' native CLE.

To register Blue Waters with the HTCondor scheduler as an OSG site, pilot jobs had to use a modified version of the GlideinWMS [13] tool. In the *Shifter* UDI based on the OSG Docker image, the tool was immersed in an OSG-like environment and, therefore, could download and execute the LIGO analysis code as usual.

A complication arose due to OSG using CVMFS [15] to distribute application codes like LIGO to the resource providers. Because CVMFS relies on FUSE [6] and the latter is not supported by the OS kernel on Blue Waters, a copy of the relevant sections of CVMFS's data hierarchy had to be stored on the Blue Waters' Lustre file system which is accessible from within the *Shifter* job.

Finally, analysis task required a data file of approximately 400 MB in size which was downloaded using GridFTP and XRootD transport protocols [19]. With GridFTP extra care was necessary not to overwhelm the data server because each GridFTP connection requires a heavy-weight runtime environment to be initialized on the data server. XRootD on the other hand is designed specifically for OSG workflows and handles multiple transfers more gracefully.

Using this setup, Blue Waters contributed approximately 8,000 node-hours to LIGO's second observation campaign, temporarily becoming the peak resource provider, and approximately 50,000 node-hours to the ATLAS project in 2017 [10].

Future HTC codes that rely on OSG resources will definitely benefit from the experiences gained and the groundwork laid by ATLAS, NANOGrav, and LIGO on Blue Waters. With the help of *Shifter*, only minimal modifications are required to enable such codes take advantage of Blue Waters, providing a new pool of compute resources otherwise unavailable to HTC codes.

QWalk, PySCF and QuantumEspresso

QWalk [18] and PySCF [14] are an electron structure and computational physics / chemistry codes. Since QWalk uses a Quantum Monte Carlo (QMC) method, it parallelizes trivially to refine its predictions using additional instances of the simulation. As such, no complex workflow manager was required and researchers were able to develop automation framework for use with *Shifter* independently.

PowerGRID

PowerGRID [4] is an MPI applications for medical magnetic resonance image reconstruction that can take advantage of GPUs. It relies on MPICH ABI compatibility to use a single executable compiled and dynamically linked with MPICH that runs under Cray's MPI stack on Blue Waters. PowerGRID employs parallelization to process multiple snapshots in parallel using MPI to farm out tasks to the cores available to the job. The per-rank code is parallelized *via* OpenACC targeting Blue Waters' NVIDIA Kepler GPUs. *Shifter* enabled the team to build a complex software stack with multiple compiler dependencies and CUDA support that they can deploy on a variety of underlying hardware.

Outlook

For all applications discussed in this paper, *Shifter* played a critical role in making their execution on Blue Waters *possible*. But why do we not see more examples like this? If we look closer at scientific applications in general, we find little consistency in the way these applications are developed. This lack of consistency leads to the use of an array of tools and packages that make the process of building applications even in a controlled environment provided by Docker very difficult. Even more so, building applications in a way that would allow them to take full advantage of the hardware provided by leadership-class computing facilities while maintaining container portability. Thus, despite all the benefits that *Shifter* brings to the world of High-Performance & Throughput Computing, there is still room for improvement.

CONCLUSIONS

We described the lessons learned and experiences gained while adopting Shifter as a container solution on the Blue Waters supercomputer. We presented a thorough and up-to-date report on its performance, functionality, issues encountered, and also the benefits and new possibilities that it enables. While some challenges remain to be solved (unsupported or chip-specific and incompatible instructions), Shifter has already provided a long-awaited solution that enabled the HPC community to run complex and atypical (for HPC) software stacks. Essentially, Shifter enabled HPC centers like Blue Waters to imitate Cloud infrastructure which is sought after by the HTC community. Over the last year, Blue Waters users have been steadily ramping up the utilization of Shifter. In addition to providing seamless access to the unique computing capabilities of Blue Waters to run HTC-tailored workflows, Shifter has provided the means to further a wave of innovation that has fused HPC and HTC resources to address grand computational challenges across science domains. We have showcased recent applications of Shifter that demonstrate the new role containers are starting to play in

maximizing the versatility and flexibility of HPC systems in accelerating scientific discovery by enabling complex and modern software stacks.

ACKNOWLEDGMENTS

This research is part of the Blue Waters sustained-petascale computing project, which is supported by the National Science Foundation (awards OCI-0725070 and ACI-1238993) and the State of Illinois. Blue Waters is a joint effort of the University of Illinois at Urbana-Champaign and its National Center for Supercomputing Applications.

We thank CERN for the very successful operation of the LHC, as well as the support staff from ATLAS institutions without whom ATLAS could not be operated efficiently. The crucial computing support from all WLCG partners is acknowledged gratefully. Major contributors of computing resources are listed in Ref [5].

The authors gladly acknowledge valuable discussions with Edgar Fajardo, Stuart Anderson, and Peter Couvares.

REFERENCES

- [1] B. P. Abbott, R. Abbott, T. D. Abbott, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. X. Adhikari, V. B. Adya, and et al. 2017. GW170817: Observation of Gravitational Waves from a Binary Neutron Star Inspiral. *Physical Review Letters* 119, 16, Article 161101 (Oct. 2017), 161101 pages. https://doi.org/10.1103/ PhysRevLett.119.161101
- [2] Maxim Belkin. 2017. Interacting with Shifter on Blue Waters. (2017). https://bluewaters.ncsa.illinois.edu/documents/10157/202012/Shifter_demo.pdf
- [3] R. S. Canon and D. Jacobsen. 2016. Shifter: Containers for HPC. (2016).
- [4] A. Cerjanic, J. L. Holtrop, G-C. Ngo, B. Leback, G. Arnold, M. Van Moer, G. LaBelle, J. A. Fessler, and B. P. Sutton. 2016. PowerGrid: An Open-Source Library for Accelerated Iterative Magnetic Resonance Image Reconstruction. In *Proc. Intl. Soc. Mag. Res. Med.* 525. http://indexsmart.mirasmart.com/ISMRM2016/PDFfiles/ 0525.html
- [5] ATLAS Collaboration. 2016. ATLAS Computing Acknowledgements 2016-2017, ATL-GEN-PUB-2016-002, 20XX. https://cds.cern.ch/record/2202407. (2016). [Online from July 2016].
- [6] Tejun Heo. 2017. The reference implementation of the Linux FUSE (Filesystem in Userspace) interface. (2017). https://github.com/libfuse/libfuse/
- [7] E. A. Huerta, Roland Haas, Edgar Fajardo, Daniel Katz, Stuart Anderson, Peter Couvares, Josh Willis, Timothy Bouvet, Jeremy Enos, William T.C. Kramer, Hon Wai Leong, and David Wheeler. 2017. BOSS-LDG: A Novel Computational Framework that Brings Together Blue Waters, Open Science Grid, Shifter and the LIGO Data Grid to Accelerate Gravitational Wave Discovery, In 2017 IEEE 13th International Conference on e-Science (e-Science). ArXiv e-prints. https: //doi.org/10.1109/eScience.2017.47
- [8] D. M. Jacobsen and R. S. Canon. 2015. Contain This, Unleashing Docker for HPC. (2015).
- [9] Andre Luckow, Mark Santcroos, Ole Weidner, Andre Merzky, Sharath Maddineni, and Shantenu Jha. 2012. Towards a Common Model for Pilot-jobs. In Proceedings of the 21st International Symposium on High-Performance Parallel and Distributed Computing (HPDC '12). ACM, New York, NY, USA, 123–124. https://doi.org/10. 1145/2287076.2287094
- [10] Mark Neubauer. 2017. Enabling Discoveries at the LHC through Advanced Computation and Machine Learning. Presented at the Blue Waters Symposium 2017. https://bluewaters.ncsa.illinois.edu/documents/10157/244350/neubauer-lhc.pdf
- [11] Mark Neubauer, Philip Chang, Rob Gardner, Dave Lesny, and Dewen Zhong. 2017. Enabling Discoveries at the Large Hadron Collider through Advanced Computation. (2017). https://bluewaters.ncsa.illinois.edu/science-teams?page= detail&psn=bafz
- [12] Ruth Pordes, Don Petravick, Bill Kramer, Doug Olson, Miron Livny, Alain Roy, Paul Avery, Kent Blackburn, Torre Wenaus, Frank Würthwein, Ian Foster, Rob Gardner, Mike Wilde, Alan Blatecky, John McGee, and Rob Quick. 2007. The open science grid. *Journal of Physics: Conference Series* 78, 1, 012057. https: //doi.org/10.1088/1742-6596/78/1/012057
- [13] Igor Sfiligoi, Daniel C Bradley, Burt Holzman, Parag Mhashilkar, Sanjay Padhi, and Frank Wurthwein. 2009. The pilot way to grid resources using glideinWMS. In Computer Science and Information Engineering, 2009 WRI World Congress on, Vol. 2. IEEE, 428–432.
- [14] Qiming Sun, Timothy C. Berkelbach, Nick S. Blunt, George H. Booth, Sheng Guo, Zhendong Li, Junzi Liu, James D. McClain, Elvira R. Sayfutyarova, Sandeep

Sharma, Sebastian Wouters, and Garnet Kin-Lic Chan. [n. d.]. PySCF: the Pythonbased simulations of chemistry framework. *Wiley Interdisciplinary Reviews: Computational Molecular Science* 8, 1 ([n. d.]), e1340. https://doi.org/10.1002/ wcms.1340

- [15] CernVM Team. 2017. CernVM File System. (2017). https://cernvm.cern.ch/ portal/filesystem
 [16] Douglas Thain, Todd Tannenbaum, and Miron Livny. 2005. Distributed computing
- [16] Douglas Thain, Todd Tannenbaum, and Miron Livny. 2005. Distributed computing in practice: the Condor experience. *Concurrency - Practice and Experience* 17, 2-4 (2005), 323–356.
- [17] S. A. Usman, A. H. Nitz, I. W. Harry, C. M. Biwer, D. A. Brown, M. Cabero, C. D. Capano, T. Dal Canton, T. Dent, S. Fairhurst, M. S. Kehl, D. Keppel, B. Krishnan, A. Lenon, A. Lundgren, A. B. Nielsen, L. P. Pekowsky, H. P. Pfeiffer, P. R. Saulson, M. West, and J. L. Willis. 2016. The PyCBC search for gravitational waves from compact binary coalescence. *Classical and Quantum Gravity* 33, 21, Article 215004 (Nov. 2016), 215004 pages. https://doi.org/10.1088/0264-9381/33/21/215004
- [18] Lucas K. Wagner, Michal Bajdich, and Lubos Mitas. 2009. QWalk: A quantum Monte Carlo program for electronic structure. J. Comput. Phys. 228, 9 (2009), 3390-3404. https://doi.org/10.1016/j.jcp.2009.01.017
- [19] D. Weitzel, B. Bockelman, D. A. Brown, P. Couvares, F. Würthwein, and E. Fajardo Hernandez. 2017. Data Access for LIGO on the OSG. ArXiv e-prints 1705.06202 [cs.DC] (May 2017).